

INTERCONNECTING NODES IN AN
OPTICAL COMMUNICATION SYSTEM

TECHNICAL FIELD OF THE INVENTION

The present invention relates generally to communication systems and, more particularly, to interconnecting nodes in an optical communication system.

DOCKET: 064731.0250

BACKGROUND OF THE INVENTION

In optical communication systems, the interconnections between nodes define, to a large extent, scalability of networks and performance of communications. A number of interconnection strategies have been proposed, with each resulting 5 topology offering various advantages and/or disadvantages compared with others. However, the ever changing requirements of communication systems often pose challenges that highlight inadequacies of existing topologies.

TOPICALS, DODGE

SUMMARY OF THE INVENTION

In accordance with the present invention, techniques for interconnecting nodes in an optical communication system are provided which substantially reduce or eliminate disadvantages and problems associated with previous techniques.

- 5 According to one embodiment of the present invention, a sub-network of an optical communication network includes a first node, a second node, a third node, and a fourth node. The first node has two input ports and two output ports, with the input ports and one of the output ports coupled to other nodes of the sub-network and the other output port operable to send signals outside of the sub-network. The second
10 node has two input ports and two output ports, with the output ports and one of the input ports coupled to other nodes of the sub-network and the other input port operable to receive signals from outside of the sub-network. The third node has two input ports and two output ports, with the output ports and one of the input ports coupled to other nodes of the sub-network and the other input port operable to receive
15 signals from outside of the sub-network. The fourth node has two input ports and two output ports, with the input ports and one of the output ports coupled to other nodes of the sub-network and the other output port operable to send signals outside of the sub-network.

Embodiments of the invention provide various technical advantages. The
20 interconnections within sub-networks and between sub-networks provide scalability. Sub-networks may be interconnected such that additional sub-networks can be added with only minimal changes to existing links. Moreover, as sub-networks are added to a network, the average path distance between nodes in the network increases linearly. That is, the average distance between nodes increases linearly as the number of nodes
25 increases. These qualities make this topology particularly attractive for relatively small communication networks, such as metropolitan area networks.

The topology provided by interconnecting nodes using these techniques also provides fault tolerance. According to particular embodiments, a communications network using this topology may continue to operate given the loss of 50 percent of
30 nodes within the network. Thus, the nodes within a network may be configured to support protection, since half of the nodes may be configured as protection nodes.

In addition, the topology proposed facilitates a relatively simple routing algorithm for communicating information between nodes. This provides, for example, a topology that supports the implementation of packet switched communications protocols. Thus, proposed topologies, while supporting existing standards, also support a migration from lightpaths to light-frames.

Other technical advantages of the present invention will be readily apparent to one skilled in the art from the following figures, descriptions, and claims. Moreover, while specific advantages have been enumerated above, various embodiments may include all, some, or none of the enumerated advantages.

EPRINTED FROM STERLING

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and its advantages, reference is now made to the following description, taken in conjunction with the accompanying drawings, in which:

5 FIGURE 1 illustrates interconnections between nodes of an optical communication system in accordance with one embodiment of the present invention;

FIGURE 2 illustrates the interconnections between the nodes using a different perspective;

10 FIGURE 3 is a block diagram illustrating exemplary functional components of a node from the communication network;

FIGURE 4 illustrates interconnections between nodes in a communication network having two sub-networks in accordance with one embodiment of the present invention;

15 FIGURE 5 illustrates an alternative perspective of the interconnections between nodes in a communication network having two sub-networks;

FIGURE 6 illustrates interconnections between nodes in a communication network having four sub-networks in accordance with one embodiment of the present invention;

20 FIGURE 7 illustrates an alternative addressing scheme for the nodes in a communication network having four sub-networks;

FIGURE 8 illustrates alternative interconnections between nodes in a communication network having four sub-networks in accordance with one embodiment of the present invention; and

25 FIGURE 9 is a flowchart illustrating a method for routing communications between nodes in the communication networks.

DETAILED DESCRIPTION OF THE DRAWINGS

FIGURE 1 illustrates an optical communication sub-network, indicated generally at 10, that includes four interconnected nodes 12. Each node 12 has two input ports and two output ports, and the solid and dotted lines illustrate the 5 interconnections of these input and output ports. As illustrated by this logical representation, the four nodes 12 forming the sub-network may be envisioned as four points of a pyramid with various interconnections forming the edges of the pyramid. This sub-network provides a building block, and one or more of these sub-networks may be interconnected to form optical communication networks of any appropriate 10 size.

In the embodiment illustrated, nodes 12 are labeled from 0 to 3, referred to as nodes n_0 , n_1 , n_2 , and n_3 . Thus, sub-network 10 includes node n_0 , node n_1 , node n_2 , and node n_3 . Each node 12 in sub-network 10 has four ports, two input ports and two output ports, and couples at least three ports to other nodes 12 within sub-network 10. 15 Thus, the solid lines represent links between nodes 12 in sub-network 10, while the dashed lines represent available links for coupling sub-network 10 to other sub-networks. Within sub-network 10, these links form a ring between nodes n_0 , n_3 , n_2 , and n_1 . Also, sub-network 10 has links across this ring from node n_1 to node n_3 and from node n_2 to node n_0 . To couple with other sub-networks, sub-network 10 uses 20 output ports at node n_0 and node n_3 and input ports at node n_1 and node n_2 . Therefore, as each node 12 within sub-network 12 has two input ports and two output ports, network 10 similarly has two input ports and two output ports. Thus, sub-network 10 may be linked with other sub-networks in order to scale for larger optical communication networks. However, in an optical communication network having 25 only four nodes 12, additional links within the sub-network may be formed. For example, the output port of node n_0 may couple to the input port of node n_2 , and the output port of node n_3 may couple to the input port of node n_1 .

FIGURE 2 is a diagram illustrating exemplary interconnections between four nodes 12 in a communication sub-network 20 in which all input ports and output ports 30 of nodes 12 are coupled to other nodes 12 within sub-network 20. Moreover, in this graphical representation, nodes 12 are illustrated in a ring configuration. This aids in illustrating the various pathways formed by the interconnections between nodes 12 of

sub-network 20. As with the lines in FIGURE 1, the solid lines in FIGURE 2 illustrate the "standard" links between nodes 12 in sub-network 20, while the dashed lines represent the ports and connections available for linking to other sub-networks or, as illustrated in this example, for linking within sub-network 20.

- 5 In operation, sub-network 20 will display favorable performance in the delivery of communications between nodes 12. From any node 12 within sub-network 20, the average hops to reach a destination is one (providing an average hop distance of unity within sub-network 20). That is, on average, a communication needs to travel only along a single link in sub-network 20 to reach its destination. For
10 example, consider a packet received by node n_1 for delivery to one of nodes 12 within sub-network 20. If the packet is destined for node n_1 , no hops are necessary to deliver the packet. For delivering the packet to node n_0 or node n_3 , only a single hop is needed, and to deliver the packet to node n_2 , two hops are needed. Thus, the average number of hops to deliver a packet from node n_1 is one. A quick examination
15 demonstrates that similar results obtain for communications from node n_0 , node n_2 , or node n_3 . Therefore, the average number of hops for a communication within sub-network 20 is one. As will be demonstrated later, as additional sub-networks are linked together, the average number of hops increases linearly. This provides particular advantages in networks having relatively small numbers of interconnected
20 nodes, such as metropolitan communication networks.

- However, while FIGURES 1 and 2 illustrate sub-networks 10 and 20 as a pyramid and a ring, respectively, it should be understood that these are merely logical representations designed to illustrate interconnections between nodes 12. Thus, the actual shapes formed by arrangements of nodes 12 within one or more sub-networks
25 may be dictated by various geographical and other criteria, and the techniques disclosed do not constrain these arrangements, but rather provide techniques for interconnecting various nodes 12. Moreover, while the links between nodes 12 are illustrated as lines, these links may include any appropriate intermediate components, such as optical amplifiers.

- 30 FIGURE 3 is a block diagram illustrating exemplary functional components of node 12 that include input ports 22, output ports 24, an optical cross-connect, add/drop multiplexer module 26, and a controller 28. In general, node 12 provides for

the receipt, routing, and communication of circuit switched and/or packet switched communications. More specifically, within a communication sub-network, three of ports 22 and 24 within node 12 will couple to other nodes 12 within the sub-network. The remaining port may couple to node 12 in another sub-network or, in certain 5 circumstances, may couple to another node 12 in the same sub-network.

Input ports 22 and output ports 24 each represent hardware, including appropriate logic, for coupling to and communicating with other communications equipment. For example, input ports 22 may include optical receivers for coupling to optical fibers and for receiving and decoding optical information signals. Similarly, 10 output ports 24 may include lasers, modulators, filters, and other appropriate equipment for coupling to optical fibers and communicating optical information signals.

Within node 12, module 26 links input ports 22 and output ports 24 to support the appropriate communication of information through node 12. Thus, module 26 15 supports the switching of lightpaths or light-frames between input ports 22 and output ports 24. In addition, module 26 couples to add/drop lines 32 and adds and removes communications from input ports 22 and output ports 24. For example, add/drop lines 32 may link to a local network, with node 12 providing the local network access to other networks.

Node 12 also includes controller 28, which controls the operation of elements 20 within node 12. For example, based on configurations and/or commands received on control line 30, controller 28 may direct the routing of communications between input ports 22 and output ports 24. Moreover, controller 28 may include routing algorithms for determining an appropriate output port 24 for delivering a communication to its 25 intended destination. In addition, while not shown in this embodiment, node 12 may include various add/drop multiplexers for introducing communications onto optical fibers or delivering communications.

However, while the particular embodiment illustrated and the preceding 30 description focus on a particular combination and arrangement of functional elements, node 12 may include any appropriate combination and arrangement of elements that include functionality for two input ports and two output ports. Thus, specific

elements illustrated and functionalities described may be separated, combined, rearranged, or otherwise modified.

- FIGURE 4 is a diagram illustrating an optical communication network 40 that includes a first sub-network 42 linked to a second sub-network 44. Network 40 provides eight nodes, from node n₀ to node n₇, with sub-network 42 formed by nodes n₀, n₁, n₂ and n₃, and sub-network 44 formed by nodes n₄, n₅, n₆ and n₇. In network 40, the two output ports from sub-network 42 link to the two input ports of sub-network 44. Thus, node n₀ links to node n₆ and node n₃ links to node n₅. Similarly, the two output ports of sub-network 44 link to the two input ports of sub-network 42. Thus, node n₄ links to node n₂ and node n₇ links to node n₁.

As evidenced by this illustration, the interconnections between nodes 12 formed using these techniques create a number of "rings" between nodes 12. These rings provide for rapid delivery of communications to appropriate destinations as well as a high tolerance against defaults. Given the interconnections shown, fifty percent of nodes 12 may fail without completely disrupting communications between the remaining nodes 12. This characteristic remains true as the number of appropriately interconnected sub-networks is increased.

- FIGURE 5 is a diagram illustrating an alternative representation of the interconnections between nodes 12 of network 40. Thus, the same interconnections between nodes 12 are illustrated as were shown in FIGURE 4, while the relative positions of nodes have been altered. As evident in this illustration, the links between nodes 12 form two counter-directional rings. In the clockwise direction, decreasingly numbered even nodes 12 are linked (with the exception of the link from node n₀ to node n₆). In the counterclockwise direction, odd numbered nodes 12 are linked in increasing order (with the exception of the link from node n₇ to node n₁). To add an additional sub-network or sub-networks into network 40, only two links need to be altered.

For example, to add a new sub-network into network 40, the link between node n₃ and node n₅ and the link between node n₄ and node n₂ may be broken and these various input and output ports then coupled to the appropriate input and output ports of the new sub-network. Similarly, the new sub-network could be added by breaking the link between node n₀ and node n₆ and the link between node n₇ and node

n₁ and then reforming appropriate links between input and output ports of the existing and additional sub-networks. This provides for a minimal disturbance of existing interconnections between nodes 12 when linking additional sub-networks. This reduces costs associated with the time, personnel, and equipment required for linking 5 and retuning links when additional nodes are added to a communication network.

FIGURE 6 illustrates an optical communication network 50 that includes four interconnected sub-networks 52, 54, 56, and 58 (referred to, respectively, as sub-network sn₀, sn₁, sn₂, and sn₃). Given the numbering system used in this illustration, each sub-network may be characterized by four nodes numbered 4n, 1+4n, 2+4n, and 10 3+4n, with n having the value of the corresponding sub-network. For example, sub-network sn₂ includes nodes n₈, n₉, n₁₀ and n₁₁, corresponding to the values obtained from 4x2, 1+4x2, 2+4x2, and 3+4x2.

Within each sub-network, nodes 4n and 3+4n serve as output ports, and nodes 1+4n and 2+4n serve as input ports. When represented in binary notation, the number 15 or address of any node 12 may be divided into a sub-network address and a local address. For example, the binary representation for node n₁₃ is 1110₂. The first two digits, 11, correspond to a sub-network address of three, while the last two digits 10 correspond to a local address of two. Thus, this address represents the second node within sub-network sn₃.

20 In the embodiment illustrated, links are formed between sub-networks according to a particular scheme. The output port at node 4n of a sub-network links to the input port of node 2+4(n-1), except for the "wrap-around" where, for example, node n₀ links to node n₁₄. The second output port of each sub-network, at node 3+4n, links to the input port of node 1+4(n+1), except for the wrap-around, for example, 25 between node n₁₅ and node n₁. This linking of ports forms two counter-rotating rings between the nodes of sub-networks within network 50.

Given the interconnections between sub-networks as illustrated in this embodiment, nodes 12 may use relatively simple algorithms for routing communications to destination nodes 12. Upon receiving a communication, node 12 30 determines whether it is destined for another node 12 within the same sub-network. For example, a communication may be addressed to a binary address having the format of xxdd. Given this address, xx identifies the destination sub-network while

- dd identifies the destination node 12 within the destination sub-network. If the destination sub-network address matches the current sub-network, node 12 routes the communication within the current sub-network. For communications within a sub-network, shortest path routing decisions are readily apparent. For example, node n_1
- 5 may route a received packet directly to node n_0 or node n_3 , or may route a packet to node n_2 via node n_3 . Thus, node n_1 , upon receiving an in sub-network communication, determines whether the destination (dd) is itself (01). If not, node n_1 determines whether the destination is 00_2 . If so, node n_1 delivers the communication to node n_0 , otherwise node n_1 delivers the communication to node n_3 . Other nodes 12
- 10 may perform similar in sub-network routing.

For communications addressed to "foreign" sub-networks, node 12 determines whether the destination sub-network is closer in a clockwise or counterclockwise direction. Node 12 then forwards the communication to and out the appropriate port of the current sub-network. For example, using binary notation, assume node 0001_2

15 receives a communication destined to node 0101_2 . The destination sub-network, 01, is closer to the current sub-network, 00_2 , in a counterclockwise direction. This indicates that the appropriate output port is 11_2 . Thus, node 0001_2 forwards the communication to the output port 11_2 of sub-network s_{n_0} , node 0011_2 , which in turn forwards the communication to node 0101_2 .

20 For sub-networks equally distant regardless of the direction traveled, the algorithm may default to a particular output port. For example, for a packet in sub-network s_{n_0} destined for sub-network s_{n_2} , the number of intervening sub-networks is equal in both a clockwise and counterclockwise direction. Thus, nodes 12 may default to a particular output port from the current sub-network given this situation.

25 This provides efficient routing that, except in certain limited circumstances, provides shortest path routing. Moreover, if shortest path routing is critical or desired, more complicated algorithms that account for path distances between specific nodes 12 may be implemented.

FIGURE 7 illustrates an optical communication network 60 that includes four

30 sub-networks 62, 64, 66, and 68. Network 60 illustrates an alternative embodiment for interconnecting sub-networks than the embodiment shown in network 50. In network 60, nodes 12 are coupled such that each node $4n$ links an output port to the

input port of node $2+4(n+1)$, and node $3+4n$ links an output port to the input port of node $1+4(n-1)$. Network 60 provides similar advantages to those provided by previous networks, however, routing algorithms may require minor alterations in order to operate appropriately.

- 5 FIGURE 8 illustrates an optical communications network 70 that includes sub-networks 72, 74, 76, and 78. As illustrated in this embodiment, when four sub-networks are interconnected, the sub-networks may be linked together in the same arrangement as nodes within each individual sub-network. That is, the output ports of sub-network sn_0 link to input ports of sub-networks sn_2 and sn_3 ; the output ports of
10 sub-network sn_1 link to input ports of sub-networks sn_0 and sn_3 ; the output ports of sub-network sn_2 link to input ports of sub-networks sn_0 and sn_1 ; and the output ports of sub-network sn_3 link to input ports of sub-networks sn_1 and sn_2 . This configuration of sub-networks may provide additional advantages when compared to other techniques for linking sub-networks. However, as with network 60, routing
15 algorithms for transmitting communications between nodes 12 in network 70 may be different than those used for other configurations of sub-networks.

- FIGURE 9 is a flowchart illustrating a method for routing communications between nodes 12. This flowchart and the following description detail the routing of a packet within a network having a configuration as detailed with respect to network
20 50. That is, within each sub-network, an output port of node $4n$ links to an input port of node $2+4(n-1)$ and an output port of node $3+4n$ links to an input port of $1+4(n+1)$.

- Upon receiving a communication, such as a packet, node 12 determines whether the destination for the communication is within the current sub-network at step 100. For example, node 12 may compare the destination sub-network address of
25 the communication with a current sub-network address. If the address is matched, node 12 routes the communication within the current sub-network at step 102. As previously discussed, node 12 may use relatively simple determinations in order to appropriately route communications to other nodes 12 within the sub-network.

- If the destination is not within the current sub-network, node 12 determines
30 whether the destination sub-network is greater than the current sub-network at step 104. If so, node 12 determines whether the difference between the destination sub-network and the current sub-network is greater than one-half the total number of sub-

networks at step 106. This permits node 12 to determine the appropriate direction for routing the communication. If the difference is greater than one-half the total number of sub-networks, the appropriate direction of travel for the communication is clockwise, and node 12 sets the output port to 00₂ at step 108. If the difference is not 5 greater than one-half the total number of sub-networks, node 12 defaults to routing the communication in a counter-clockwise direction and thus sets the output port to 01₂ at step 110. After determining the appropriate output port, node 12 routes the communication to the selected output port within the sub-network at step 112.

If the destination sub-network is not greater than the current sub-network, the 10 logic for determining an appropriate output port inverts. Thus, given the circumstance, node 12 determines whether the difference between the current sub-network and the destination sub-network is less than one-half the total of sub-networks at step 114. If so, the appropriate direction of travel for the communication is clockwise, and node 12 sets the output port to 00₂ at step 116. Otherwise, node 12 15 defaults to forwarding the communication in a counterclockwise direction and sets the output port to 01₂ at step 118. After appropriately setting the output port, node 12 routes the communication to the selected output port within the sub-network at step 112.

The preceding flowchart illustrates only an exemplary method of operation, 20 and nodes 12 may use any suitable techniques for routing communications between elements of a network. Thus, many of the steps in this flowchart may take place simultaneously and/or in different orders than as shown. In addition, nodes 12 may use methods with additional steps, fewer steps, and/or different steps, so long as the 25 methods remain appropriate. For example, given various techniques for linking sub-networks within a network, the methods used for routing communications should be tailored accordingly. Therefore, the example provided by this flowchart is merely illustrative, and node 12 contemplates using any suitable steps for routing communications.

Although the present invention has been described in several embodiments, a 30 myriad of changes and modifications may be suggested to one skilled in the art, and it is intended that the present invention encompass such changes and modifications as fall within the scope of the present appended claims.